arXiv:1612.03236v1 [cs.CV] 10 Dec 2016

Co-localization with Category-Consistent CNN Features and Geodesic Distance Co-Propagation

Hieu Le *[1], Chen-Ping Yu[3], Gregory Zelinsky[1,2] and Dimitris Samaras[1]

[1]Department of Computer Science, Stony Brook University
[2]Department of Psychology, Stony Brook University
[3]Department of Psychology, Harvard University

Abstract

Co-localization is the problem of localizing categorical objects using only positive sets of example images, without any form of further supervision. This is a challenging task as there is no pixel-level annotations. Motivated by human visual learning, we find the common features of an object category from convolutional kernels of a pre-trained convolutional neural network (CNN). We call these category-consistent CNN features. Then, we co-propagate their activated spatial regions using superpixel geodesic distances for localization. In our first set of experiments, we show that the proposed method achieves state-of-the-art performance on three related benchmarks: PASCAL 2007, PASCAL-2012, and the Object Discovery dataset. We also show that our method is able to detect and localize truly unseen categories, using six held-out ImagNet subset of categories with state-of-the-art accuracies. Our intuitive approach achieves this success without any region proposals or object detectors, and can be based on a CNN that was pre-trained purely on image classification tasks without further fine-tuning.

1. Introduction

Convolutional neural networks (CNNs) have been widely applied in the general problem of object localization and detection. The task is to detect a target object's location and its spatial coverage in an image in the form of bounding boxes [30, 6, 25, 24]. To localize objects from images, typically a model is given images of category exemplars for training. Critically, these training samples have precise object-level annotations, such as segmentations or bounding boxes. The models can be fine-tuned from a pre-trained network and utilize region proposals for candidate object locations [39, 14, 40, 9, 4], or trained end-to-end [25, 35, 24]. These models have demonstrated high performance in localizing objects from learned categories, and further fine-tuning is required in order to accommodate novel object

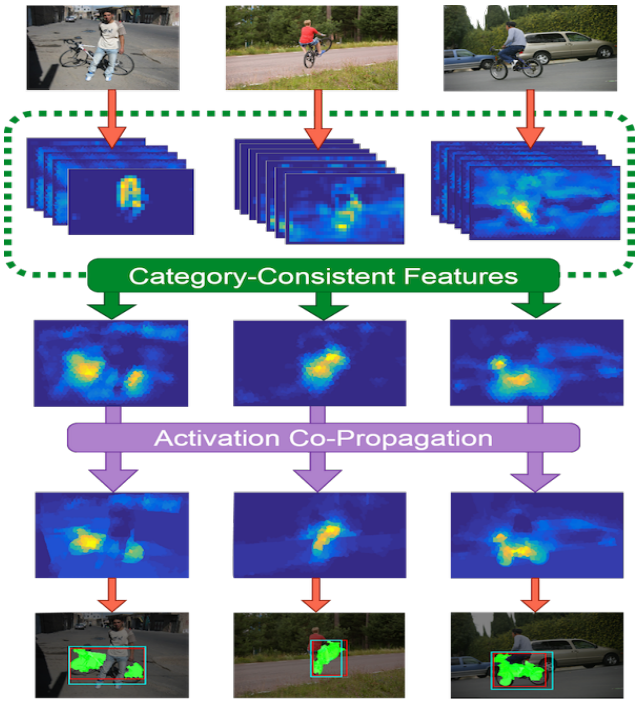

Figure 1. **Object co-localization with CCFs and geodesic distance co-propagation**. From a set of images containing a common object, first we find the CCFs - the group of features that consistently have high responses to the object images of the same class. The CCFs then are used to form an activation map for each image, followed by geodesic distance co-propagation to highlight the exact regions of the objects.

*hle@cs.stonybrook.edu

categories [22] .

Co-localization is the more challenging problem of localizing objects from only the set of positive image examples of the category without any object-level annotations. The lack of negative examples and detailed annotations hinders the use of supervised methods for the co-localization task. Recent methods typically utilize existing region proposal methods for generating a number of candidate regions for objects and object parts, followed by matching or selecting the region with the highest confidence score [18, 34, 7, 21]. However, region and object proposals are part of a research problem of its own, and have drawbacks such as lack of repeatability, reduced detection performance with a large number of proposals, and lead to difficult balance in precision and recall[15].

Our method, however, does not require any object proposals or object detectors to perform co-localization. The main idea in our work is that objects of the same class share common features or parts. Moreover, these commonalities are central to both, the category representation and the detection and localization of the object. By finding those object categorical features, their joint locations can act as a single-shot object detector. This idea is also grounded in human visual learning, where it is suggested that people detect common features from examples of the category, as part of the object-learning process [38]. We do this by obtaining the CNN features of the provided set of positive images, in order to select the ones that are highly and consistently activated, which we denote as Category-Consistent CNN Features (CCFs) We then use these CCFs to discover the rough object locations, and demonstrate an effective way to co-propagate the feature activations into a stable object for precise co-localization. Figure 1 illustrates the pipeline of our proposed framework.

In more detail, our approach begins with a CNN that has been pre-trained for image classification on ImageNet. Then, the images of the target category are passed through the network. We identify the last-layer convolutional filters that have highly and consistently activated feature maps as the CCFs. The CCFs' feature maps are combined into a single normalized activation probability map, where the highly activated region directly implies the rough object location, since the CCFs represent object parts or object-associated features. The CCF step allows us to bypass the need for region proposals. Then, the activation map is partitioned into superpixels and weighted by the superpixel geodesic distance into an object-likelihood map such that the responses of the object-associated features propagate over the region of the entire object. Finally, the precise object location can be obtained by placing a tight bounding box around the thresholded object-likelihood map.

The three main contributions of this work are: **1.** We propose a novel CCF extraction method that can automatically highlight the rough initial object regions, which acts as a single-shot detector. **2.** We introduce an effective method of feature co-propagation for generating a stable object region using superpixel geodesic distances on the original images. **3.** Our method achieves state-of-the-art performance for object co-localization on the VOC 2007 and 2012 datasets [11], the Object Discovery dataset [28], and the six held-out ImageNet subset categories. Furthermore, our framework is fully unsupervised, objects are discovered using just positive image exemplars. We are able to accurately localize objects without needing any region proposals.

2. Related work

Co-localization is related to work on weakly supervised object localization (WSOL) [33, 31, 8, 36, 5, 26, 37] since both share the same objective: to localize objects from an image. However, since WSOL allows the use of negative examples, designing the objective function to discover the information of the object-of-interest is less challenging: WSOL-based methods achieve higher performance on the same datasets as compared to co-localization methods, due to the allowed supervised training. For instance, Wang et al. [36] uses image labels to evaluate the discrimination of discovered categories in order to localize the objects. Ren et al. [26] adopts a discriminative multiple instance learning scheme to compensate the lack of object-level annotations to localize the objects based on the most discriminative instances. Because of the supervision that is required by those methods, it is not trivial for WSOL approaches to be directly applied to the co-localization scenarios.

One challenge of co-localization is to define the criteria for discovering the objects without any negative examples. To fill the gap, state-of-the-art co-localization methods such as [21, 34, 7, 18] employ object proposals as part of their object discovery and co-localization pipelines. Tang et al. [34] use the measure of objectness [2] to generate multiple bounding boxes for each image, followed by an objective function to simultaneously optimize the image-level labels and box-level labels. Such settings allow the use of discriminative cost function [16]. This is also used in the work of co-localization on video frames [18]. Cho et al. [7] also starts from object proposals, their method shares the same spirit with the deformable part model [12] where the objects are discovered and localized by matching common object parts. Most recently, Li et al. [21] study the confidence score distribution of a supervised object detector over the set of object proposals to define an objective function, that learns a common object detector with similar confidence score distribution. All the aforementioned methods heavily depend on the quality of object proposals.

Our work approaches the problem from a different perspective. Instead of trying to fill in the gap of the negative data and annotations that are unavailable, we find the com-

mon features shared by the objects from the positive images. Then, we use the joint locations of those features as our single-shot object detector. This allows us to bypass the need for utilizing a region or object proposal algorithm as a fist step. Our subsequent step refines the detected object features into a stable object by co-propagating their activations together. We describe the details of our 2-step approach in the following sections.

3. Extracting Category-Consistent CNN Features

Our proposed method consists of two main steps. The first step is to find the CCFs of a category, and obtain their combined feature map that contains aggregated CCF activations over the rough object region. Then, the CCF activations are co-propagated into a stable object using superpixel geodesic distances on the original images.

Given a set of n object images from the same class and a CNN that has been pre-trained to contain sufficient visual features, we first compute the m feature maps from the k last-layer convolutional kernels over the n images. Then, we obtain an $m \times n$ activation matrix with each row being the activation vector a of a kernel containing the maximum values of the kernel's feature maps.

Specifically, for each kernel: $A_{i,j} = \max(F(i,j))$, where $F(i,j)$ is the feature map of kernel i, given image j that has been forward-passed through the CNN. The activation matrix A therefore describes the max-response distributions of all kernels to all category images.

Our goal in this step is to identify a subset of representative kernels from the global set of candidate kernels, that contain common features from the positive images of the same class. This implies that the activation vectors of the kernels should have high values over all vector elements, since there is at least one instance of the object on every image. Conceptually, the kernels that we seek correspond to object parts, or some object-associated features. To find the CCFs, we compute the pair-wise similarities between all pairs of kernels' activation vectors, and cluster them using k-means. The kernels from the cluster with the highest mean activation correspond to the CCFs. The similarity $s_{i,j}$ between two CNN kernels i and j, can be defined as the L_p distance between two activation vectors A_i and A_j.

The sets of positive images are only required for the identification of the CCF kernels. The CCF kernels can then be used to generate the rough object location in an image in a single-shot: given an image from the target category, the feature maps of the CCFs are combined to form a single activation map. Since each CCF corresponds to an object part of object-associate features, the densely activated area of the activation map indicates the rough location of the target object. The final activation map is normalized into a probability map that sums to 1. In Figure 2 we show the identified CCFs for the bus category, where the activated regions describe bus-related features and all fall within the spatial extent of the objects.

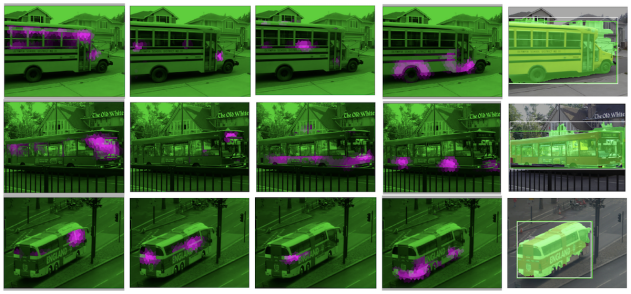

Figure 2. **Examples of our CCFs for *bus* category.** Each row is a different example image, and each column is the activation (in violet) of a single CNN feature in the set of our CCFs. The last column shows the final co-localization results.

4. Stable Object Completion via Co-Propagating CCF Activations

The activation probability map from the CCFs automatically points out only the rough location of the object. It does not ensure a reliable object localization due to: **1.**The higher layer of a CNN does not guarantee a kernel's receptive-field size to cover the area of an entire object. **2.**While the feature maps contain spatial information, they have unprecise object locations due to previous max-pooling layers. **3.**The CNN was trained discriminatively. Hence, only discriminative features of each object may be localized rather than the the whole object.

In order to obtain the complete region that corresponds to the object, we compute geodesic distances between superpixels on the original image. In essence, the geodesic distance compactly encodes the similarity relationship between the two superpixels' image contents. The similarity is computed via object boundary detection algorithm [20]. Therefore, the smaller geodesic distance between the two superpixels, the more likely that they belong to the same object. Based on this characteristic, we propose a simple and effective method to highlight the object region from the activation probability map, that is both low-resolution and contains non-smooth feature activations.

Given an input image, we oversegment it into superpixels. The geodesic distance $d_{i,j}$ between a pair of superpixels i, j is computed based on the graph built from the boundary probability map, which is done similarly as the method proposed by Krähenbühl and Koltun [20]. We take the combined activation map that was obtained in the CCF identification step, and assign an energy value to each superpixel by averaging its corresponding pixel values found

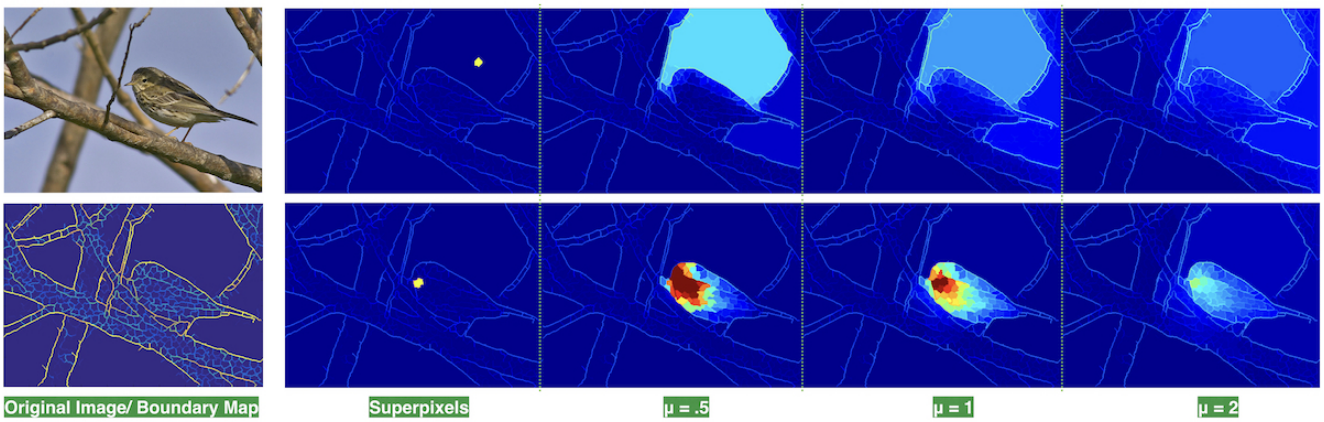

Figure 3. **Geodesic distance propagation.** The first column is the original image and the boundary map extracted by [23]. The second column marks the position and initial activation of two selected superpixels, for illustrating the effect of our activation propagation for a single superpixel. The top superpixel is located in the background while the bottom one in located in the region of the category object (*bird*). The next columns illustrate the propagated activations from the two superpixels to all other superpixels, with the controlling parameter μ set at 0.5,1 and 2, respectively.

in the activation map. For k superpixels, we denote the resulting flattened $k \times 1$ superpixel activation vector as $\mathbf{E}$. Vector $\mathbf{E}$ can be considered as the initial likelihood of each superpixel being within the object.

Next, we perform the geodesic distance propagation to localize the object. The main idea is that if two superpixels are likely to belong to the same object, then they should have similar geodesic distances and activations. This concept has been similarly adopted by various works in terms of interactive image segmentation and matting [3, 13], and we find it to suit specially well for our purpose. Therefore, we seek to obtain a global activation map that has regions of highly boosted or co-propagated activations by superpixels of similar geodesic distances, mediated by some level of consistency by formulating this co-propagating mechanism into a $k \times k$ co-propagation matrix $\mathbf{W}$, such that $W_{i,j}$ is the normalized amount of co-propagation between superpixel i and j, with a parameter μ for controlling the amount of activation diffusion:

$$W_{i,j} = \frac{\exp(-d_{i,j} \times \mu^{-1})}{\sum_{k=1}^{N} \exp(-d_{i,k} \times \mu^{-1})}. \quad (1)$$

Finally, we apply W to the activation vector E directly:

$$E' = \mathbf{WE}, \quad (2)$$

where E' is the co-propagated activation vector of the image, containing the globally boosted activations of the superpixels based on their pair-wise geodesic distances to all other superpixels. This allows us to fill in each superpixel on the image with their respective values from E', and normalize the co-propagated superpixel map by dividing every pixel by the max value of the map. The result is an object-likelihood map, on which we apply a global threshold to obtain the region as our final object co-localization result. Finally, a tight bounding box is placed around the maximum coverage of the thresholded regions within an image.

Figure 3 illustrates the effects of activation propagation for two selected superpixels. The top row corresponds to the case of a background superpixel. This superpixel has small geodesic distances to a large set of other superpixels. Hence, the activation values propagated from this superpixel to others are relatively small. In the second row, we consider a superpixel corresponding to the bird. In this case, high activation values are only propagated from this superpixel to other superpixels that also correspond to the bird. Hence, we filter out the undesirable high activation of the background regions surrounding the object, while also balancing the activation of all superpixels that reside within the same object. The first column of figure 3 shows the original image and its boundary map. The second column marks the position and initial activation of the two selected superpixels. The next columns illustrate the activation propagating from the selected superpixel to all other superpixels with a varying degree of controlling parameter μ set at 0.5,1 and 2, where warmer values indicate higher activations. It can be seen from the third column that since the background superpixel has small geodesic distances to many other superpixels, the activations being propagated from this superpixel to others are equally small; in contrast, the activations propagated from the bottom superpixel mostly fall into its close neighboring superpixels and ones that belong to the object's region. As μ increases, the amount of propagation is more evenly and widely spread, but with a lower overall magnitude.

Methods	Airplane	Car	Horse	Mean
[19]	21.95	0.00	16.13	12.69
[16]	32.93	66.29	54.84	51.35
[17]	57.32	64.04	52.69	58.02
[27]	74.39	87.64	63.44	75.16
[18]	71.95	93.26	64.52	76.58
[7]	82.93	**94.38**	75.27	84.19
Ours	**84.15**	**94.38**	**78.49**	**85.67**

Table 4. **Experiment on Object Discovery Dataset.** Highest performances are labeled in bold.

5. Experiments

We evaluate our proposed 2-step framework with different parameter settings to illustrate different characteristics of our method. We also evaluate our method on multiple benchmarks, with intermediate and final results to show the localization effects of our proposed method. In all of our experiments, we used the last convolutional layer of a VGG-19 network [32] that was pre-trained on ImageNet [29] as our CCF kernel pool. We used $k = 5$ for kmeans clustering in the CCF identification step. The geodesic distances are computed using the Structured Forests soft boundary [23]. The control parameter μ for activation co-propagation was set at 0.5. The final global threshold for obtaining the object region from the object-likelihood map was set at 0.25 for all images.

5.1. Evaluation metric and datasets

We use the conventional CorLoc metric [10] to evaluate our co-localization results. The metric measures the percentage of images that contain correctly localized results. An image is considered correctly localized if there is at least one ground truth bounding box of the object-of-interest having more than 50% Intersection-over-Union (IoU) score with the predicted bounding box. To benchmark our method performance, we evaluate our method on three commonly used datasets for the problem of co-localization. These are VOC 2007 and 2012 [11], and Object Discovery dataset [28]. For experiments on VOC datasets, we followed previous works [7, 18, 21] that used all images on the *trainval* set excluding the images that only contain the object instances annotated as *difficult* or *truncated*. For experiments on the Object Discovery dataset, we used the 100-image subset following [27] in order to make an appropriate comparison with related methods. The ground truth bounding box for each image in the Object Discovery dataset is defined as the smallest bounding box covering all the segmentation ground truth of the object.

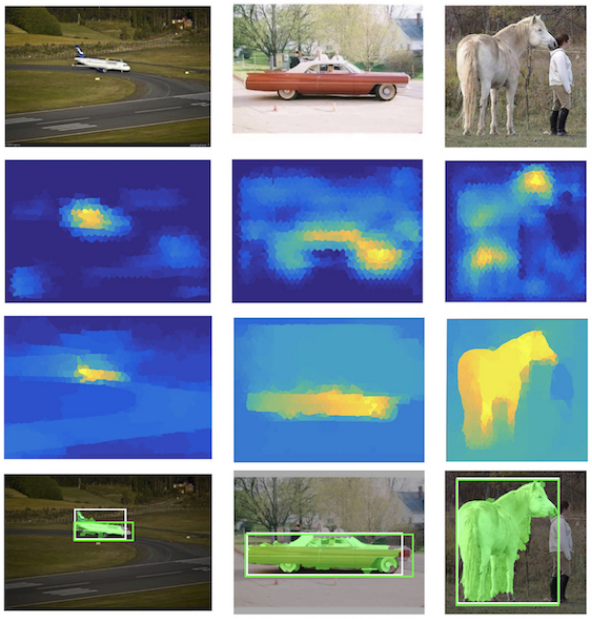

Figure 4. **Object co-localization results on Object Discovery dataset.** From the top row to the bottom row: input image, combined activation map from the identified CCFs, propagated object-likelihood map, and resulting bounding boxes. We also depict in green the pixels with the object-region that's predicted by our method. Our predicted bounding boxes are colored as white while the ground truth bounding boxes are green.

5.2. Comparison to state-of-the-art co-localization methods

We first evaluate our method on the 100-image subset of Object Discovery dataset which contains objects of three classes, namely *airplane*, *car*, and *Horse*. There are 18, 11, and 7 noisy images in each class, respectively. Table 4 reports the co-localization performance of our approach in comparison with the state-of-the-art methods on image co-localization [16, 27, 17, 18, 7, 21]. In this small scale setting, our method outperformed other methods in both in individual object classes and overall.

Three examples of our co-localization approach on the Object Discovery dataset are illustrated in figure 4. The second row shows the combined activation map from the set of identified CCFs, that acted as our single-shot object detector. It is apparent that the combined activation maps already provided object estimates that were quite accurate to the location of the actual object in the images, with different parts of each object getting high values such as the tail of the airplane, the wheel of the car, or head and tail of the horse. All these values were co-propagated based on the superpixel geodesic distances, resulted in the images shown in the third row. The co-propagated object-likelihood maps on the third row show that the sporadic activations on

VOC	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean
[18]	32.8	17.3	20.9	18.2	4.5	26.9	32.7	41.0	5.8	29.1	**34.5**	31.6	26.1	40.4	17.9	11.8	25.0	27.5	35.6	12.1	24.6
[7]	50.3	42.8	30.0	18.5	4.0	62.3	**64.5**	42.5	8.6	**49.0**	12.2	44.0	64.1	57.2	15.3	9.4	30.9	**34.0**	61.6	**31.5**	36.6
[21]	**73.1**	45.0	43.4	27.7	6.8	53.3	58.3	45.0	6.2	48.0	14.3	47.3	**69.4**	66.8	**24.3**	12.8	**51.5**	25.5	65.2	16.8	40.0
ours	56.3	**46.75**	**43.8**	**29.6**	**8.7**	**62.8**	54.6	**74.9**	**8.8**	42.1	27.7	**50.5**	59.2	**71.0**	19.3	**13.9**	45.4	31.0	**68.9**	9.6	**41.2**

Table 1. **CorLoc scores of our approach and state-of-the-art co-localization methods on Pascal VOC 2007 dataset.**

VOC	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean
[7]	57.00	41.20	36.00	26.90	5.00	**81.10**	**54.60**	50.90	**18.20**	54.00	**31.20**	44.90	61.80	48.00	13.00	11.70	51.40	**45.30**	64.60	**39.20**	41.80
[21]	**65.70**	57.80	47.90	28.90	6.00	74.90	48.40	48.40	14.60	**54.40**	23.90	50.20	**69.90**	68.40	**24.00**	14.20	52.70	30.90	**72.40**	21.60	43.80
ours	63.06	**67.79**	**50.39**	**41.06**	**19.73**	79.20	44.50	**74.79**	15.27	40.39	27.42	**68.40**	69.36	**74.34**	21.17	**14.29**	**53.37**	39.69	69.41	15.46	**47.45**

Table 2. **CorLoc scores of our approach and state-of-the-art co-localization methods on Pascal VOC 2012 dataset.**

VOC	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean
[33]	42.40	46.50	18.20	8.80	2.90	40.90	73.20	44.80	5.40	30.50	19.00	34.00	48.80	65.30	8.20	9.40	16.70	32.30	54.80	5.50	30.38
[31]	67.30	54.40	34.30	17.80	1.30	46.60	60.70	68.90	2.50	32.40	16.20	58.90	51.50	64.60	18.20	3.10	20.90	34.70	63.40	5.90	36.18
[8]	56.60	58.30	28.40	20.70	6.80	54.90	69.10	20.80	9.20	50.50	10.20	29.00	58.00	64.90	36.70	18.70	56.50	13.20	54.90	59.40	38.84
[37]	37.70	58.80	39.00	4.70	4.00	48.40	70.00	63.70	9.00	54.20	**33.30**	37.40	**61.60**	57.60	30.10	31.70	32.40	**52.80**	49.00	27.80	40.16
[5]	66.40	59.30	42.70	20.40	21.30	**63.40**	**74.30**	59.60	21.10	58.20	14.00	38.50	49.50	60.00	19.80	39.20	41.70	30.10	50.20	44.10	43.69
[26]	79.20	56.90	46.00	12.20	15.70	58.40	71.40	48.60	7.20	**69.90**	16.70	47.40	44.20	**75.50**	**41.20**	**39.60**	47.40	32.20	49.80	18.60	43.91
[36]	**80.10**	**63.90**	**51.50**	14.90	**21.00**	55.70	74.20	43.50	**26.20**	53.40	16.30	56.70	58.30	69.50	14.10	38.30	**58.80**	47.20	49.10	**60.90**	**47.68**
ours	56.3	46.75	43.8	**29.6**	8.7	62.8	54.6	**74.9**	8.8	42.1	27.7	50.5	59.2	71.0	19.3	13.9	45.4	31.0	**68.9**	9.6	41.2

Table 3. **CorLoc scores of our approach and state-of-the-art weekly-supervised-object-localization methods on Pascal VOC 2007 dataset.**

the background have been smoothed out evenly, and that the non-smooth object parts and their associated activations have been boosted and completed into complete and stable objects with significantly higher activation magnitudes. This shows that our two-step framework was able to generate informative single-shot object detection using the CCFs, and the subsequent stable object region via activation co-propagation.

The PASCAL VOC 2007 and 2012 datasetes both contains realistic images of 20 object classes with significantly larger numbers of images per class. These datasets are more challenging than the Object Discovery dataset due to the diversity of viewpoints, and the complexity of the objects. Table 1 reports our performance on the VOC 2007 dataset. Our approach outperforms the state-of-the-art method [21] by 1.2% on average, and acquires the highest scores for 11 out of 20 classes. Results on VOC 2012 dataset, that has twice the number of images than the VOC 2007 dataset, are shown in table 2. Our method also achieves significantly better results than state-of-the-art methods with an 3.65% increase on average, acquiring the highest scores for 10 out of 20 classes.

We also compare our method with state-of-the-art weekly-supervised object localization methods [33, 31, 8, 36, 5, 26, 37], which is summarized in table 3. While our results did not achieve the overall state-of-the-art, our approach was able to outperform the methods in 3 of the 20 classes, without using any negative examples or any form of supervised learning.

5.3. Category-consistent CNN features selection analysis

In this section, we provide an analysis to justify our CCF selection method, where we conducted additional experiments with the same configurations but using different subsets of CNN features for the initial object detection step. After clustering the last-layer CNN kernels based on their image-level activations, these clusters were sorted in a decreasing order by the clusters' average activations. We then obtained the rough object locations using individual clusters and thresholded on those maps directly to obtain the object locations. Their respective average CorLoc scores on the VOC 2007 and 2012 dataset are reported in table 5.

The table shows that the co-localization performance followed in the exact order of the clusters based on their average activations, suggesting that the most representative features were indeed members of the top cluster. The visualized examples are shown in figure 5, with an image from the *dog* and *motorbike* category, respectively. For each image, the first row is the results of our method when using the first cluster (ranked by their average activation) and the second row shows the results of our method when using the third

Dataset	1st	2nd	3rd	4th	5th
VOC07	41.24	38.85	31.73	29.62	23.12
VOC12	47.45	42.35	37.70	33.91	25.76

Table 5. **Co-localization performance of our method on the VOC 2007 and 2012 dataset.** Each column indicates which top cluster was taken as the CCF cluster (out of 5 total clusters), and the corresponding average CorLoc score (%) by using the selected cluster of features for co-localization.

Cluster	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv
1st	6.3	4.5	4.1	3.7	8.0	7.6	4.1	3.7	22.9	6.8	8.4	5.7	4.1	5.3	16.2	9.2	5.7	20.1	10.4	7.6
2nd	18.6	16.2	8.8	19.7	29.5	17.4	20.5	17.4	38.3	12.7	36.5	6.1	17.8	23.6	13.5	24.0	5.7	11.5	22.1	19.7
3rd	30.7	9.6	13.5	28.3	10.0	27.5	32.4	23.6	7.2	24.4	8.0	22.9	13.5	30.7	26.0	24.0	20.9	32.4	26.2	30.3
4th	29.5	43.0	40.6	24.6	37.5	28.7	28.7	23.6	20.1	34.4	32.8	38.1	35.2	29.7	30.5	28.5	38.3	26.0	25.2	28.5
5th	15.0	26.8	33.0	23.6	15.0	18.8	14.3	31.6	11.5	21.7	14.3	27.3	29.5	10.7	13.9	14.3	29.5	10.0	16.2	13.9

Table 6. **The number of kernels (%) per cluster.** The percentage of kernels in each cluster based on the activation vectors described in Section 3. The clusters were formed by using K-mean with $k = 5$, and the kernels were taken from the last convolutional layer of VGG-19, on each class of the VOC 2012 dataset.

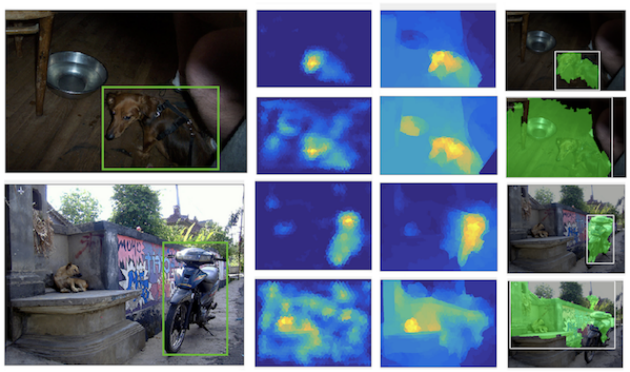

Figure 5. **Two examples illustrating the effect of our feature selection method.** For each image, the first row is the results of our method when using the first cluster and the second row is the results of our method when using the third cluster. The green bounding box is the ground truth and the predicted bounding boxes are colored as white, the predicted object regions are masked as green.

cluster. It is clear that the combined activation maps from the third cluster failed to detect and estimate the object locations, and ultimately lead to incorrect object localization results. This indicates that the selection of the top cluster is essential, and the CCFs could not be chosen arbitrarily.

We furthermore validate our feature selection method by evaluating the performance of our approach when the CCFs were identified from more than one cluster, and the results are shown in table 7. This experiment shows that large amount of kernels do not provide enough object specificity, and therefore resulting in a similar performance decline as in table 5, such that performance decreases when more clusters were added.

Dataset	1	1-2	1-3	1-4	1-5
VOC07	41.24	41.00	38.22	34.80	33.25
VOC12	47.45	46.68	44.14	40.65	39.33

Table 7. **Co-localization performance of our methods on the VOC 2007 and 2012 dataset.** Each column indicates how many top clusters were taken as the CCF clusters (out of 5 total clusters), and the corresponding average corLoc score (%) by using the selected clusters of features for co-localization. For example, the last column indicates that all available features were used for co-localization.

We also report the number of kernels per cluster in table 6. Looking at table 6 and table 2 together, they suggest that there are multiple classes that require just a small amount of kernels in order to be localized decently. For example, the *cat* class used less than 20 kernels (3.7% of the 512 total kernels) to achieve the state-of-the-art CorLoc score of 74.9% in VOC 2012 dataset. In general, the table shows that the first cluster only contains a small number of kernels, and the results suggest that these small number of specific CCFs are sufficient for the co-localization performed on the three benchmark datasets.

5.4. Geodesic distance co-propagation analysis

Geodesic distance acts as a refinement step in our pipeline to get rid of the background as well as boosting the activation within the object region. To evaluate the effect of geodesic distance co-propagation, we simply compared the performances of our method on VOC 2007 and 2012 datasets with and without geodesic distance co-propagation, and the results are reported in table 8. The results show that geodesic distance co-propagation significantly improved the co-localization accuracy by more than 6% in absolute CorLoc score for both dataset, which means it is an important process subsequent to our initial CCF object detection step.

Dataset	Without GDP	With GDP
VOC07	34.45	41.24
VOC12	40.97	47.45

Table 8. **Co-localization performance of our methods on VOC 2007 and 2012 dataset with and without using geodesic distance propagation.** Using geodesic distance co-propagation increased the average CorLoc score (%) of our approach by 6.79% in CorLoc for VOC 2007 dataset, and 6.48% in CorLoc for VOC 2012 dataset.

5.5. Unseen classes from ImageNet subset

Our VGG-19 model was pre-trained on ImageNet's 1000 classes. While VOC 2007 and 2012 are different datasets from ImageNet, there are significant overlaps between the object categories in VOC and ImageNet. For example, the "motorbike" class of VOC datasets is equivalent to the "moped" class of ILSVRC 2012 dataset.

ImageNet	chipmunk	rhino	stoat	racoon	rake	wheelchair
# of images (Full)	307	213	237	1301	466	420
# of images (Smaller from [21])	159	89	159	103	146	174

Table 9. Number of images in each of the 6 subsets selected from ImageNet dataset, collected by us directly from the ImageNet website (Full), and the smaller set by Li et al.[21] (Smaller from [21]).

ImageNet	chipmunk	rhino	stoat	racoon	rake	wheelchair	mean
[7]	26.60	81.80	44.20	30.10	8.30	35.30	37.72
[21]	44.00	81.80	67.30	41.80	**14.50**	39.30	48.12
Ours	**72.33**	**88.76**	**75.59**	**83.16**	11.64	**49.53**	**63.50**

Table 10. CorLoc scores [10] (%) of our approach and state-of-the-art co-localization methods on the 6 smaller subsets of ImageNet collected by Li et al. [21].

ImageNet	chipmunk	rhino	stoat	racoon	rake	wheelchair	mean
Ours	76.87	87.79	80.59	74.10	50.64	54.05	70.67

Table 11. CorLoc scores [10] (%) of our approach and state-of-the-art co-localization methods on the 6 full subsets of ImageNet.

Table 10 reports the co-localization performance of our method on the 6 smaller subsets in comparison with the two state-of-the-art methods. The table shows that our method outperformed [21] and [7] by a large margin on average, and all but one individual classes.

The six subsets of the ImageNet dataset, chosen by Li et al. [21], are held-out categories from the 1000-label classification task, which means they do not overlap with the 1000 classes used to train VGG-19. We show that our method is generalizable to truly novel object categories with the six held-out ImageNet subset classes. The images and the corresponding bounding box annotations were downloaded from ImageNet website [1], and Table 9 shows the numbers of images in each class with available bounding box annotations when we accessed the ImageNet dataset website. It is noticeable that paper [21] used an even smaller set with less images (table 9). To compare with the methods of [21] and [7], we first conduct our experiment on the same smaller sets of images that was used in [21]. Then, We test our method on the full set of ImageNet held-out categories that we downloaded from the ImageNet dataset.

Table 10 reports the co-localization of our method compares to [21] on the smaller ImageNet subset, and table 11 reports the performance of our method on the ImageNet full subsets. The results show that our method significantly outperforms the competing methods in the smaller subset, with even higher accuracies for the full subset in Table 11. This result demonstrates that our method is robust to detect and localize truly unseen categories using previously learned CNN features.

5.6. Qualitative results

We show some examples of our co-localization results in figure 7 and 8. The results show that the bounding boxes generated by our proposed framework accurately match the ground truth bounding boxes. It is apparent that our results generate well-covered object regions, that has the potential to well delineate the objects in majority of the cases. The figures also show that the objects were able to be accurately co-localized with various sizes and locations.

Figure 6. **Some failed examples of our approach.** While these examples did not have significant coverage with the ground truth boudning box in terms of CorLoc, the mistakes were still accurate, in the sense that the areas were detected but not the spatial extent.

Figure 6 illustrates three failure scenarios of our approach. While these three examples did not cover the ground truth bounding box sufficiently, but they were not far off. Some analysis suggests that these failures were due to some CCFs that are shared by multiple categories, and that the object boundaries may not have been strong enough (i.e. bottle and boat).

6. Conclusion

In this work, we proposed a fully unsupervised 2-step approach for the problem of co-localization, that uses only positive images, and without any region or object proposals. Our method is motivated by human vision, people implicitly detect the common features of category examples to learn the representation of the class. We show that the identified category-consistent features can also act as an effective first-pass object detector. This idea is implemented

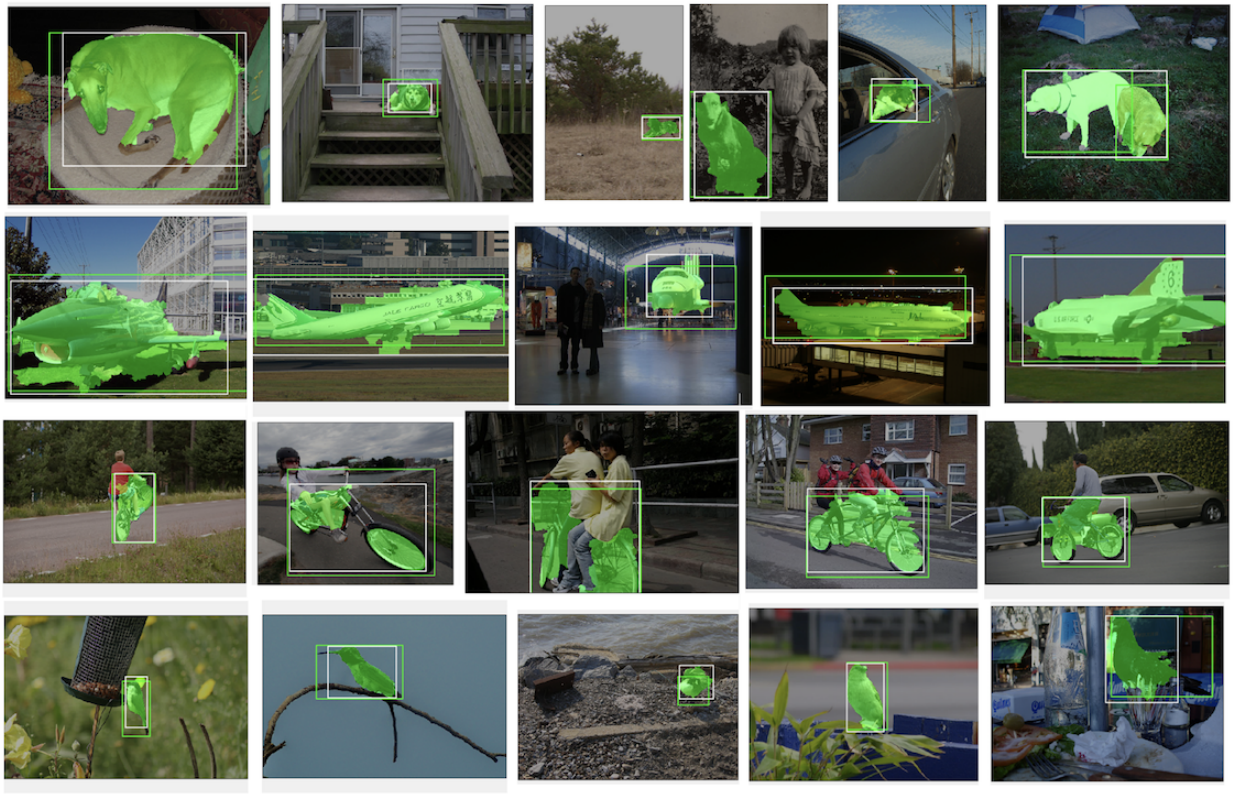

Figure 7. **Some example results of object co-localization with our CCFs and geodesic distance co-propagation.** The results show the bounding boxes that we generate match the ground truth bounding boxes very well, even when objects are not located centrally in the image. In addition, our co-localized object pixel-level regions (pixels colored in green) show well delineated shape in most cases.

Figure 8. **Some example results of our object co-localization on class "chipmunk"** from the six held-out ImageNet subset categories. The results show co-localized object pixel-level regions (pixels colored in green), the bounding boxes from our method (white), and ground truth bounding boxes (green).

by finding the group of CNN features that are highly and consistently activated by a given positive set of images. The result of this first step generates a rough but reliable object location, and acts as a single-shot object detector. Then, we aggregate the activations of the identified CCFs, and co-propagate their activations so that the activations over the true object region are boosted, while the activations over the background region are smoothed out. This effective activation refinement step allowed us to obtain accurately co-localized objects in terms of the standard CorLoc score with bounding boxes. We achieved new state-of-the-art performance on the three commonly used benchmarks. In the

future, we plan to extend our method to generate unsupervised object co-segmentations.

References

[1] Imagenet. http://image-net.org/. Accessed: 2016-11-15.

[2] Bogdan Alexe, Thomas Deselaers, and Vittorio Ferrari. Measuring the objectness of image windows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(11):2189–2202, November 2012. ISSN 0162-8828. doi: 10.1109/TPAMI.2012.28. URL http://dx.doi.org/10.1109/TPAMI.2012.28.

[3] X. Bai and G. Sapiro. Geodesic matting: a framework for fast interactive image and video segmentation and matting. *IJCV*, 2008.

[4] A. Bency, H. Kwon, H. Lee, S. Karthikeyan, and B. Manjunath. Weakly supervised localization using deep feature maps. In *ECCV*, 2016.

[5] H. Bilen, M. Pedersoli, and T. Tuytelaars. Weakly supervised object detection with convex clustering. In *CVPR*, 2015.

[6] J. Caicedo and S. Lazebnik. Active object localization with deep reinforcement learning. In *ICCV*, 2015.

[7] M. Cho, S. Kwak, C. Schmid, and J. Ponce. Unsupervised object discovery and localization in the wild: part-based matching with bottom-up region proposals. In *CVPR*, 2015.

[8] Ramazan Gokberk Cinbis, Jakob Verbeek, and Cordelia Schmid. Multi-fold MIL Training for Weakly Supervised Object Localization. In *CVPR 2014 - IEEE Conference on Computer Vision & Pattern Recognition*, Columbus, United States, June 2014. IEEE. URL https://hal.inria.fr/hal-00975746.

[9] Ramazan Gokberk Cinbis, Jakob Verbeek, and Cordelia Schmid. Weakly supervised object localization with multi-fold multiple instance learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016.

[10] Thomas Deselaers, Bogdan Alexe, and Vittorio Ferrari. Weakly supervised localization and learning with generic knowledge. *International Journal of Computer Vision*, 100 (3):275–293, 2012. ISSN 1573-1405. doi: 10.1007/s11263-012-0538-3. URL http://dx.doi.org/10.1007/s11263-012-0538-3.

[11] M. Everingham, S. Eslami, L. Gool, C. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *IJCV*, 2015.

[12] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645, 2010.

[13] L. Grady. Random walks for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2006.

[14] K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. In *ECCV*, 2014.

[15] J. Hosang, R. Benenson, P. Dollar, and B. Schiele. What makes for effective detection proposals. *arXiv*, 2015.

[16] A. Joulin, F. Bach, and J. Ponce. Discriminative clustering for image co-segmentation. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.

[17] A. Joulin, F. Bach, and J. Ponce. Multi-class cosegmentation. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.

[18] A. Joulin, K. Tang, and F.-F. Li. Efficient image and video co-localization with frank-wolfe algorithm. In *ECCV*, 2014.

[19] Gunhee Kim, E. P. Xing, Li Fei-Fei, and T. Kanade. Distributed cosegmentation via submodular optimization on anisotropic diffusion. In *2011 International Conference on Computer Vision*, pages 169–176, Nov 2011. doi: 10.1109/ICCV.2011.6126239.

[20] Philipp Krähenbühl and Vladlen Koltun. *Geodesic Object Proposals*, pages 725–739. Springer International Publishing, Cham, 2014. ISBN 978-3-319-10602-1. doi: 10.1007/978-3-319-10602-1_47. URL http://dx.doi.org/10.1007/978-3-319-10602-1_47.

[21] Y. Li, L. Liu, C. Shen, and A. van den Hengel. Image co-localization by mimicking a good detector's confidence score distribution. In *ECCV*, 2016.

[22] M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Is object localization for free? weakly-supervised learning with convolutional neural networks. In *CVPR*, 2015.

[23] Larry Zitnick Piotr Dollar. Structured forests for fast edge detection. In *ICCV*. International Conference on Computer Vision, December 2013. URL https://www.microsoft.com/en-us/research/publication/structured-forests-for-fast-edge-detection/.

[24] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016.

[25] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.

[26] Weiqiang Ren, Kaiqi Huang, Dacheng Tao, and Tieniu Tan. Weakly supervised large scale object localization with multiple instance learning and bag splitting. *IEEE*

Trans. Pattern Anal. Mach. Intell., 38(2):405–416, February 2016. ISSN 0162-8828. doi: 10.1109/TPAMI.2015.2456908. URL http://dx.doi.org/10.1109/TPAMI.2015.2456908.

[27] M. Rubinstein, A. Joulin, J. Kopf, and C. Liu. Unsupervised joint object discovery and segmentation in internet images. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.

[28] M. Rubinstein, A. Joulin, J. Kopf, and C. Liu. Unsupervised joint object discovery and segmentation in internet images. In *CVPR*, 2013.

[29] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

[30] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y LeCun. OverFeat:integrated recognition, localization and detection using convolutional networks. In *ICLR*, 2014.

[31] Zhiyuan Shi, Timothy M. Hospedales, and Tao Xiang. Bayesian joint topic modelling for weakly supervised object localisation. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2013.

[32] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.

[33] Parthipan Siva and Tao Xiang. Weakly supervised object detector learning with model drift detection. In *Proceedings of the 2011 International Conference on Computer Vision*, ICCV '11, pages 343–350, Washington, DC, USA, 2011. IEEE Computer Society. ISBN 978-1-4577-1101-5. doi: 10.1109/ICCV.2011.6126261. URL http://dx.doi.org/10.1109/ICCV.2011.6126261.

[34] K. Tang, A. Joulin, L.-J. Li, and F.-F. Li. Co-localization in real-world images. In *CVPR*, 2014.

[35] J. Tompson, R. Goroshin, A. Jain, Y. LeCun, and C. Bregler. Efficient object localization using convolutional networks. In *CVPR*, 2015.

[36] Chong Wang, Weiqiang Ren, Kaiqi Huang, and Tieniu Tan. Weakly supervised object localization with latent category learning. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part VI*, pages 431–445, 2014.

[37] Xinggang Wang, Zhuotun Zhu, Cong Yao, and Xiang Bai. Relaxed multiple-instance svm with application to object discovery. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, ICCV '15, pages 1224–1232, Washington, DC, USA, 2015. IEEE Computer Society. ISBN 978-1-4673-8391-2. doi: 10.1109/ICCV.2015.145. URL http://dx.doi.org/10.1109/ICCV.2015.145.

[38] C.-P. Yu, J. Maxfield, and G. Zelinsky. Searching for category-consistent features: a computational approach to understanding visual category representation. *Psychological Science*, 2016.

[39] N. Zhang, J. Donahue, R. Girshick, and T. Darrell. Part-based r-cnns for fine-grained category detection. In *ECCV*, 2014.

[40] Y. Zhang, K. Sohn, R. Villegas, G. Pan, and H. Lee. Improving object detection with deep convolutional networks via bayesian optimization and structured prediction. In *CVPR*, 2015.